arXiv:1401.3325v2 [cs.LO] 18 Nov 2015

Equilibria in multi-player multi-outcome infinite sequential games

Stéphane Le Roux
Département d'informatique
Université libre de Bruxelles,* Belgium
Stephane.Le.Roux@ulb.ac.be

Arno Pauly
Département d'informatique
Université libre de Bruxelles, Belgium
&
Computer Laboratory
University of Cambridge, United Kingdom
Arno.Pauly@cl.cam.ac.uk

We investigate the existence of certain types of equilibria (Nash, ε-Nash, subgame perfect, ε-subgame perfect, Pareto-optimal) in multi-player multi-outcome infinite sequential games. We use two fundamental approach: one requires strong topological restrictions on the games, but produces very strong existence results. The other merely requires some very basic determinacy properties to still obtain some existence results. Both results are transfer results: starting with the existence of some equilibria for a small class of games, they allow us to conclude the existence of some type of equilibria for a larger class.

To make the abstract results more concrete, we investigate as a special case infinite sequential games with real-valued payoff functions. Depending on the class of payoff functions (continuous, upper semi-continuous, Borel) and whether the game is zero-sum, we obtain varying existence results for equilibria.

In the context of infinite-duration games played on finite graphs we provide similar results showing that if the two-player win/lose versions are finite memory determined, this can be transferred to the general case under some conditions.

Our results hold for games with two or up to countably many players.

1 Introduction

The present article continues the research programme to investigate sequential games in a very general setting, which was initiated by the first author in [19, 21]. It extends the conference paper [22]. This programme reunites two mostly separate developments in the study of games. On the one hand, the first development is the investigation of variations on solution concepts for games, and of different formalizations of the preferences of the players, which primarily happened inside game theory proper. Related to that, game theory has also seen an interest in relaxing the continuity and convexity assumptions of Nash's original existence theorem [27], both for an example and references to further work see [28]. So typically, the countably infinite is absent from game theory: sets are either finite, or in cases such as randomized strategies, have the structure of the continuum.

On the other hand, the study of infinite sequential games has a long history in logic. Many variations on the rules of games have been studied, albeit mostly restricted to zero-sum games with two players and two outcomes. Thus, here the continuum is entirely absent (disregarding

*A significant part of this work was completed while the first author was a postdoctoral researcher at the University of Darmstadt, Germany.

its internal occurrence in the set of potential plays), and the countably infinite is only used for time, not for e.g. the number of agents or outcomes.

In our work we study infinite sequential games with perfect information (i.e. generalized Gale-Stewart games [11]) in a setting as general as possible. We can have any countable number of players, and we investigate various ways to represent preferences. Some results do put restrictions on the number of distinguished outcomes (as being countable), and some require a (generalized) zero-sum condition. A similar synthesis of the approaches is found in [41, 12, 9, 35].

A classical result by Martin [24] established that such a game played by two players who only care about whether or not the play falls into some fixed Borel set is determined, i.e. admits a winning strategy for one of the players. While determinacy can be understood as a special case of existence of Nash equilibrium, Nash equilibrium is often regarded as an unsatisfactory solution concept for sequential games in game theory. Subgame perfect equilibria are a more convincing solution concept from a rationality perspective. This article studies these two concepts, and as best responses are not always available, we also investigate the existence of ε-Nash equilibria and ε-subgame perfect equilibria. Furthermore, we also study Pareto-optimal equilibria.

The proofs that we provide fall into two broad categories: some of our results are obtained by lifting Borel determinacy to more complicated settings, similar to [19, 21] or to a sketched observation by Mertens and Neyman in [26]. Other results are based on topological arguments to show strong existence results, albeit at the cost of continuity requirements in the game characterizations. Both proof techniques provide general results in rather abstract settings. In this sense, our main results are Theorem 8 on the one hand, and Theorems 23, 24, 27 (which share parts of their proofs) on the other hand.

A common theme of the results is that they are transfer principles: they tell us how to take a pre-existing result on a more restricted class of games and obtain from it a result for a more general class of games. Theorem 8 is then used to transfer the existence of subgame perfect equilibria from finite sequential games to certain infinite ones. With Theorems 23, 24, 27, we can in particular extend Borel determinacy to yield equilibria in multi-player multi-outcome settings.

In an attempt to make the rather abstract results somewhat more accessible, we shall consider in addition the corollaries obtained in the situation where the goals of the players are to maximize real-valued payoff functions. Distinguishing properties here are continuity, upper or lower semicontinuity and Borel measurability. These settings have been studied before [10, 26, 35, 9], and usually our corollaries improve upon known results by extending them from the case of finitely many players to the case of countably many players. An overview of past and new results is given in the table on Page 3.

Our results showcase which requirements are actually needed for which aspects of determinacy, and as such may contribute to the understanding of strategic behaviour in general.

Additionally, the emergence of quantitative objectives in addition to qualitative structure in traditional verification/synthesis games [2], such as mean-payoff parity games [5, 4], provides an area of applications for abstract theorems about the existence of equilibria. Existence results in such settings are usually not trivial, but are proven together with the introduction of the setting – thus we do not answer open questions, but are hopeful that our results may be useful in the future. In these cases, strategies realized by finite automata are particularly relevant, and we provide a dedicated transfer result for these in Section 7.

Having results for countably many players is important for applications if multi-agent interactions in open systems are studied. In order to employ an equilibrium existence result for

Overview

Payoff functions	Type	ε-Nash	Nash	ε-subgame perfect	subgame perfect
continuous	finitely many players	yes -	yes -	yes -	yes [10, Corollary 4.2]
continuous	countably many players	yes -	yes -	yes -	yes Corollary 10
upper semi-continuous	finitely many players	yes -	yes Corollary 29	yes [35, Theorem 2.1]	no Example 43
upper semi-continuous	countably many players	yes -	yes Corollary 29	?	no Example 43
lower semi-continuous	finitely many players	yes -	no Example 16	yes [9, Theorem 2.3]	no -
lower semi-continuous	countably many players	yes -	no Example 16	no [9, Subsection 4.3]	no -
Borel	zero-sum	yes -	no Example 16	yes Corollary 26	no -
Borel	finitely many players	yes Corollary 15	no Example 16	no Example 43	no -
Borel	countably many players	yes Corollary 15	no Example 16	no Example 43	no -

Note: the table shows for any relevant combination of properties of the payoff functions, type of the game and type of equilibria whether such equilibria exist always or whether there is a counterexample. The results as listed here pertain to having finitely many choices at each stage of the game. There is no difference between two-player games and games with finitely many players in any situation we investigate. As we consider perfect information games only, any zero-sum game is understood to be a two-player game. The combination of semi-continuity and zero-sum would imply continuity, and is thus left out. For continuous payoff functions, we already see complete positive results without the zero-sum condition, and thus do not mention it explicitly either. Both Corollary 15 and Example 16 seem to be folklore results.

finitely many players, a bound on the number of agents involved in the interaction might need to be common knowledge from the beginning on. Our results on the other hand easily enable a setting where additional agents may join the interaction later on, and only the number of agents who have acted in the past is finite.

2 Background

In our most abstract definition, a game is a tuple $\langle A, (S_a)_{a\in A}, (\prec_a)_{a\in A}\rangle$ consisting of a non-empty set A of *agents* or *players*, for each agent $a \in A$ a non-empty set S_a of *strategies*, and for each agent $a \in A$ a *preference* relation $\prec_a \subseteq \left(\prod_{a\in A} S_a\right) \times \left(\prod_{a\in A} S_a\right)$. The generic setting suffices to introduce the notion of a Nash equilibrium: a *strategy profile* $\sigma \in \left(\prod_{a\in A} S_a\right)$ is called a Nash equilibrium, if for any agent $a \in A$ and any strategy $s_a \in S_a$ we find $\neg\left(\sigma \prec_a \sigma_{a\mapsto s_a}\right)$, where $\sigma_{a\mapsto s_a}$ is defined by $\sigma_{a\mapsto s_a}(b) = \sigma(b)$ for $b \in A\setminus\{a\}$ and $\sigma_{a\mapsto s_a}(a) = s_a$. In words, no agent prefers over a Nash equilibrium some other situation that only differs in her choice of strategy.

We will give additional structure to games in two primary ways: in Section 3 we add topologies to the strategy spaces, and then impose some topological constraints on both strategy

spaces and preferences. Beyond that, we will consider games where strategy spaces and preferences are derived objects from more structured variants of games. One such variant is the infinite sequential game:

Definition 1 (Infinite sequential game, cf. [19, Definition 1.1])**.** An *infinite sequential game* is an object
$\langle A, C, d, O, v, (\prec_a)_{a\in A}\rangle$ complying with the following.

1. A is a non-empty set (of agents).
2. C is a non-empty set (of choices).
3. $d : C^* \to A$ (assigns a decision maker to each stage of the game).
4. O is a non-empty set (of possible outcomes of the game).
5. $v : C^\omega \to O$ (assigns outcomes to infinite sequences of choices).
6. Each $\prec_a$ is a binary relation over O (modeling the preference of agent a).

The intuition behind the definition is that agents take turns to make a choice. Whose turn it is depends on the past choices via the function d. Over time, the agents thus jointly generate some infinite sequence, which is mapped by v to the outcome of the game. Note that using a single set of actions C for each step just simplifies the notation, a generalization to varying action sets is straightforward.

The infinite sequential games are linked to abstract games as follows: the agents remain the agents and the strategies of agent a are the functions $s_a : d^{-1}(\{a\}) \to C$. We can then safely regard a strategy profile as a function $\sigma : C^* \to C$ whose *induced play* is defined below, where for an infinite sequence $p \in C^\omega$ we let p_n be its n-th value, and $p_{\leq n} = p_{<n+1} \in C^*$ be its finite prefix of length n.

Definition 2 (Induced play and outcome, cf. [19, Definition 1.3])**.** Let $s : C^* \to C$ be a strategy profile. The *play* $p = p^\gamma(s) \in C^\omega$ *induced by* s *starting at* $\gamma \in C^*$ is defined inductively through its prefixes: $p_n = \gamma_n$ for $n \leq |\gamma|$ and $p_n := s(p_{<n})$ for $n > |\gamma|$. Also, $v \circ p^\gamma(s)$ is the *outcome induced by* s *starting at* γ. The play (resp. outcome) induced by s is the play (resp. outcome) induced by s starting at ε.

In the usual way to regard an infinite sequential game as a special abstract game, an agent prefers a strategy profile σ to σ', iff he prefers the outcome induced by σ to the outcome induced by σ'. And indeed we shall call a strategy profile of an infinite sequential game a Nash equilibrium, iff it is a Nash equilibrium with these preferences. In a certain notation overload, we will in particular use the same symbols for the preferences over strategy profiles and the preferences over outcomes.

However, there is a certain criticism of this choice as being not rational: essentially, the resulting concept of a Nash equilibrium means that players can use empty threats – declarations they would play in a certain way from a position onwards, even if that would be against their own interests once the position is reached, as long as this threat keeps other players from moving to that position. This can be fixed by considering subgame perfect equilibria [38]. We understand these as the Nash equilibria derived from a different translation of preferences from infinite sequential games to abstract games. (Similar remarks were made in [17, Lemma 144 in Section 7.2.3, Section 7.3.2].)

Definition 3. Given an infinite sequential game $\langle A, C, d, O, v, (\prec_a)_{a\in A}\rangle$, let the subgame perfect preferences[1] $\prec_a^{sgp} \subseteq C^{C^*} \times C^{C^*}$ be defined by $\sigma \prec_a^{sgp} \sigma'$ iff $\exists \gamma \in C^*$ such that $p^\gamma(\sigma) \prec_a p^\gamma(\sigma)$.

The subgame perfect equilibria of $\langle A, C, d, O, v, (\prec_a)_{a\in A}\rangle$ are the Nash equilibria of $\langle A, (C^{d^{-1}(\{a\})})_{a\in A}, (\prec_a^{sgp})_{a\in A}\rangle$.

We consider a further variant, namely the *infinite sequential games with real-valued payoffs*, which can (but do not have to) be understood as a special case of infinite sequential games.

Definition 4. An infinite sequential game with real-valued payoffs is a tuple $\langle A, C, d, (f_a)_{a\in\mathbb{N}}\rangle$ where A, C, d are as above, and $f_a : C^\omega \to [0, 1]$ is the payoff function of player a.

Such a game can be identified with the infinite sequential game $\langle A, C, d, [0, 1]^A, v, (\prec_a)_{a\in A}\rangle$ where $v(p) = (f_a(p))_{a\in A}$ and for $x, y \in [0, 1]^A$, we set $x \prec_a y$ iff $x_a < y_a$.

As with the introduction of subgame perfect equilibria, we can consider infinite sequential games with real-valued payoffs as infinite sequential games in a different way, which then gives rise to another commonly studied equilibrium concept, namely ε-Nash equilibria. Depending on how we then translate from infinite sequential games to abstract games, we obtain also ε-subgame perfect equilibria. Given some $\varepsilon > 0$, we define the relation $\prec_a^\varepsilon \subseteq [0, 1]^A \times [0, 1]^A$ by $x \prec_a^\varepsilon y$ iff $y_a - x_a > \varepsilon$. Using $\prec_a^\varepsilon$ in place of $\prec_a$ in Definition 4 then provides the above-mentioned equilibrium notions.

For infinite sequential games with real-valued payoffs, every Nash equilibrium (w.r.t. the standard preferences) is an ε-Nash equilibrium; and every subgame perfect equilibrium is an ε-subgame perfect equilibrium. For infinite sequential games, every subgame perfect equilibrium is a Nash equilibrium, in particular, any ε-subgame perfect equilibrium is an ε-Nash equilibrium.

We use *antagonistic game* to refer to two-player games with preferences satisfying $\prec_a = \prec_b^{-1}$, where $x \prec^{-1} y \Leftrightarrow y \prec x$.

We proceed to recall a few more notions that are only tangentially related to the formulation of our results, but that do show up in the proofs.

Definition 5. A two-player win-loose game is a tuple $\langle C, D, W\rangle$ with $D \subseteq C^*$ and $W \subseteq C^\omega$. It corresponds to the infinite sequential game $\langle\{a, b\}, C, d, \{0, 1\}, v, \{<, <^{-1}\}\rangle$ where d is defined via $d^{-1}(\{a\}) = D$ and v is defined via $v^{-1}(\{1\}) = W$.

Finally, we will extend the notion of the induced play. Given some partial function $s :\subseteq C^* \to C$, we define the consistency set $P(s) \subseteq C^\omega$ by:

$$P(s) = \{p(\sigma) \mid \sigma : C^* \to C \wedge \sigma|_{\mathrm{dom}(s)} = s\}$$

Pareto-optimality

Pareto-optimality provides a notion of social desirability in game theory, and can be used both to pick particularly *nice* equilibria, and to investigate whether the strategic interaction is costly in some sense[2]. The fundamental idea is that a Pareto-optimal outcome cannot be improved to everyone's satisfaction. Pareto-optimality is often only defined for linear preferences (or, slightly more general, strict weak orders), and its extension to general preferences is not obvious. The two natural choices are:

[1] Note that the translation of preferences in the following definition does not preserve acyclicity. Preservation could be ensured, *e.g.*, by giving the nodes a linear "priority" order, in a lexicographic fashion. This, however, would complicate the definition against little benefit for the point that we want to make.

[2] Similar to the (quantitative) *price of stability*, see [36].

Definition 6. An outcome is *realizable* in some game, if it is assigned to some sequence of choices. We call an outcome o *Pareto-optimal*, if there is no other realizable outcome q such that for some player a we find $o \prec_a q$ and for no player b we have $q \prec_a o$. We call an outcome o *weakly Pareto-optimal*, if there is no other realizable outcome q such that for some player a we find $o \prec_a q$ and for all players b we have that $o \preceq_b q$.

Note that for linear preferences, both notions coincide. We shall call a Nash equilibrium (weakly) Pareto-optimal iff it induces a (weakly) Pareto-optimal outcome.

Games played on finite graphs

This subsection shall provide the fundamental notions for games played on finite graphs used for Section 7. A two-player win/lose game played on a finite graph is specified by a directed graph (V, E) where every vertex has an outgoing edge, a starting vertex $v_0 \in V$, two sets $V_0 \subseteq V$ and $V_1 := V \setminus V_0$, and a *winning condition* $W \subseteq V^\omega$. Starting from v_0, the players move a token along the graph, ω times, with player $a \in \{0, 1\}$ picking and following an outgoing edge whenever the current vertex lies in V_a. Player 0 wins iff the infinite sequence of visited vertices is in W.

For $a \in \{0, 1\}$ let H_a be the set of finite paths in (V, E) starting at v_0 and ending in V_a. Let $H := H_0 \cup H_1$ be the possible *finite histories* of the game, and let $[H]$ be the infinite ones. For clarity we may write $[H_g]$ instead of $[H]$ for a game g. A *strategy* of player $a \in \{0, 1\}$ is a function of type $H_a \to V$ such that $(v, s(hv)) \in E$ for all $hv \in H_a$. A pair of strategies (s_0, s_1) for the two players induces a run $p \in V^\omega$ defined as follows: let $s = s_0 \cup s_1$ and set $p(0) = v_0$ and $p(n + 1) = s(p(0)p(1) \dots p(n))$. For all strategies s_a of player a let $H(s_a)$ be the finite histories in H that are compatible with s_a, and let $[H(s_a)]$ be the infinite ones. A strategy s_a is said to be winning if $[H(s_a)] \subseteq W$, *i.e.* a wins regardless of her opponent's moves.

A strategy is finite memory, if it can be computed by a finite automaton. More specifically, a strategy s_a with memory size m is a strategy that can be implemented via a function $\sigma : V \times \{0, 1\}^m \to V \times \{0, 1\}^m$, where $(v, M) \mapsto \pi_2 \circ \sigma(v, M)$ describes the memory update if the memory content was M just before reaching the vertex v, and $(v, M) \mapsto \pi_1 \circ \sigma(v, M)$ describes the strategical choices when the memory content is M at the vertex v, which is only relevant when $v \in V_a$. Up to isomorphism we can set the memory content of our strategies to be 0^m at the beginning of the game: $M_\sigma(v_0) := 0^m$, and then $M_\sigma(hv) := \pi_2 \circ \sigma(v, M_\sigma(h))$ for all $hv \in H$, and $s_a(hv) := \pi_1 \circ \sigma(v, M_\sigma(hv))$ for all $hv \in H_a$.

A (general) game played on a finite graph is specified by a directed graph (V, E), a set of agents A, a cover $\{V_a\}_{a \in A}$ of V via pairwise disjoint sets, the starting vertex v_0, and for each player a a preference relation $\prec_a \subseteq [H] \times [H]$. The notions of strategies and induced plays generalize in the obvious way. In particular, instead of a pair of strategies (one per player), we consider families $(s_a)_{a \in A}$, which are called strategy profiles. The notion of Nash equilibrium applies as before.

There is a minor formal discrepancy between the our definitions of game played on a finite graph and infinite sequential game: in the infinite sequential case, the set of available actions is the same at each time, while the number and identity of available outgoing edges in the graph may of course vary with the current vertex. This can be fixed easily by fixing some surjection from C to each set of outgoing edges. In order to keep notation simple, we tacitly ignore this aspect.

The transfer of results from the two-player win/lose case to the general case relies on the idea that each general game induces a collection of two-player win/lose games:

Definition 7 (Future game and one-vs-all threshold game)**.**
Let $g = \langle (V, E), v_0, A, \{V_a\}_{a\in A}, (\prec_a)_{a\in A}\rangle$ be a game played on a finite graph.

- For all $h \in H$ ending with some v_1, let $g_h := \langle (V, E), v_1, A, \{V_a\}_{a\in A}, (\prec_a^h)_{a\in A}\rangle$ be called the future game of g after h, where for all $hp, hp' \in [H]$ we have $v_1 p \prec_a^h v_1 p'$ iff $hp \prec_a hp'$.
- Let $a_0 \in A$ and $p \in [H]$, the one-vs-all threshold game $g^{a_0,p}$ for a_0 and p is the win-lose two-player game played on (V, E), starting at v_0, with vertex subsets V_{a_0} and $\bigcup_{a\in A\setminus\{a_0\}} V_a$, and with winning set $\{p' \in [H] \mid p \prec_{a_0} p'\}$ for the first player.

We will employ some additional restrictions on preferences: a preference relation $\prec \subseteq [H] \times [H]$ is called *prefix-linear*, if $p \prec q \Leftrightarrow hp \prec hq$ for all $p, q, hp \in [H]$. It is *prefix-independent*, if $p \prec q \Leftrightarrow hp \prec q$ and $q \prec p \Leftrightarrow q \prec hp$ for all $p, q, hp \in [H]$. Clearly, a prefix-independent preference is prefix-linear.

As a further generalization, we will consider *automatic-piecewise prefix-linear* preferences $\prec$. Here, there is an equivalence relation on H that is decidable by a finite automaton, such that if h' is in $\overline{h}$, the equivalence class of h, then for all $hp, hq, h'p, h'q \in [H]$ we have $hp \prec hq \Leftrightarrow h'p \prec h'q$. Since the graph is finite, it is possible to refine the equivalence relation while preserving its finiteness properties and such that two histories ending with different vertices belong to different classes, which we will assume without loss of generality.

3 The continuity argument

A strong transfer result can be obtained using topological arguments alone, with the reasoning being particularly well-adapted to a formulation in synthetic topology (originally [7], [31] for a short introduction). Consider games in normal form, with potentially countably many agents with strategy spaces $\mathbf{S}_1, \mathbf{S}_2, \ldots$. Our first condition is that each $\mathbf{S}_i$ be compact (subsequently, by Tychonoff's Theorem, also $\Pi_{i\in\mathbb{N}}\mathbf{S}_i$). This restriction is very common and usually combined with continuity of the outcome function, as it avoids pathological games such as *pick-the-largest-natural-number*. Our second condition is that each preference relation $\prec_i$ is open (as a subset of $(\Pi_{i\in\mathbb{N}}\mathbf{S}_i) \times (\Pi_{i\in\mathbb{N}}\mathbf{S}_i)$).[3] In the reading of synthetic topology, this means that any agent will be able to eventually confirm that he prefers a given strategy profile to another, provided he does indeed do so. We shall call a class of games $\mathcal{G}$ satisfying these conditions (in a uniform way) to be compact-strategies, open-preferences. Uniformity here means that we assume a topology on $\mathcal{G}$ such that the function mapping a game to the preferences is continuous itself.

We will write $\mathcal{O}(\mathbf{X})$ for the hyperspace of open subsets of $\mathbf{X}$, and $\mathcal{K}(\mathbf{X})$ for the hyperspace of compact sets. By $\mathcal{C}(\mathbf{X}, \mathbf{Y})$ we denote the space of continuous functions from $\mathbf{X}$ to $\mathbf{Y}$, in particular $\mathcal{C}(\mathbb{N}, \mathbf{X})$ we denote the space of sequences in $\mathbf{X}$. For precise definitions, see [31]. There we also find that the following operations are continuous:

[3] There actually is a third condition, that any strategy space is overt. A space $\mathbf{X}$ is overt, if $\{\emptyset\} \subseteq \mathcal{O}(\mathbf{X})$ is a closed set, i.e. if there is a way to detect non-emptiness of open subsets. This condition could only ever fail in a constructive reading, but is always valid for topological spaces in classical logic. Synthetic topology however would also allow us to read *continuous map* to mean *computable map*, in which case overtness becomes non-trivial. In this reading, though, we actually obtain an algorithmic result.

1. [4] $\exists : \mathcal{O}(\mathbf{X} \times \mathbf{Y}) \to \mathcal{O}(\mathbf{X})$, defined by $\exists(U) = \{x \in \mathbf{X} \mid \exists y \in \mathbf{Y}\ (x, y) \in U\}$
2. $\bigcup : \mathcal{C}(\mathbb{N}, \mathcal{O}(\mathbf{X})) \to \mathcal{O}(\mathbf{X})$
3. $^C : \mathcal{O}(\mathbf{X}) \to \mathcal{K}(\mathbf{X})$, provided that $\mathbf{X}$ is compact.
4. NonEmptyValue : $\mathcal{C}(\mathbf{X}, \mathcal{K}(\mathbf{Y})) \to \mathcal{O}(\mathbf{X})$ defined by NonEmptyValue$(f) = \{x \in \mathbf{X} \mid f(x) \neq \emptyset\}$.

Theorem 8. Let $\mathcal{G}$ be compact-strategies, open-preferences, and let $\mathcal{G}' \subseteq \mathcal{G}$ be a dense subclass. If every $G \in \mathcal{G}'$ has a Nash equilibrium, then every $G \in \mathcal{G}$ has a Nash equilibrium.

Proof. By combining our continuous operations, we may obtain the set of all games in $\mathcal{G}$ with a Nash equilibrium as an open set in the following way:

$$\mathcal{NE} := \text{NonEmptyValue}\left(G \mapsto \left(\bigcup_{i \in \mathbb{N}} \exists(\prec_i^G)\right)^C\right) \in \mathcal{O}(\mathcal{G})$$

Formulating the individual steps in words: with $\prec_i$ being open[5], we immediately obtain that the set of all strategy profiles such that player i has a better response is uniformly open in the game. Taking the union over all players again yields an open set, which complement now is the closed set NE(G) of all Nash equilibria of the respective game G. As this is a subset of the compact space $(\Pi_{i \in \mathbb{N}} \mathbf{S}_i)$, we can even treat NE$(G)$ as a compact set uniformly in G. By the synthetic definition of compactness, we obtain that $\{G \in \mathcal{G} \mid \text{NE}(G) = \emptyset\} \subseteq \mathcal{G}$ is an open set.

Because we have assumed $\mathcal{G}'$ to be dense in $\mathcal{G}$, we see that if any game in $\mathcal{G}$ would fail to have a Nash equilibrium, this would imply that some game in $\mathcal{G}'$ would fail, too, contrary to the assumption. □

Lemma 9. Consider sequential games with continuous payoff-functions and finite choices sets. We find:

1. The subgame-perfect preferences produce a compact-strategies, open-preference class $\mathcal{S}$.
2. The games with payoffs fully determined after finitely many moves are a dense subset $\mathcal{S}^f \subseteq \mathcal{S}$.
3. All games in $\mathcal{S}^f$ have a Nash equilibrium.

Proof.
1. First, note that the mapping $(\gamma, \sigma) \mapsto p_\gamma(\sigma)$ is continuous, so for any fixed $\gamma \in C^*$, $\{(\sigma, \sigma') \mid f_a(p_\gamma(\sigma)) < f_a(p_\gamma(\sigma'))\}$ is an open set. By taking countable union, we learn that $\prec_a^{sgp}$ is open. Compactness and overtness of the strategy spaces are straightforward.
2. As the argument in (1) is uniform in the continuous functions f_a, it suffices to argue that the payoff functions $f : C^\omega \to [0, 1]$ depending only on some finite prefix of the input are dense in $\mathcal{C}(C^\omega, [0, 1])$. A countable base for the applicable topology is found in all

$$\{f \mid \forall \gamma \in C^k, p \in C^\omega\ .\ p_{\leq k} = \gamma \Rightarrow f(p) \in (x_\gamma, y_\gamma)\}$$

[4] In general, the continuity of this map would require $\mathbf{Y}$ to be overt. As explained above, in a classical reading, this condition is always satisfied.

[5] And $\mathbf{S}_i$ being overt, see above.

for

$$k \in \mathbb{N}, x_{(\cdot)}, y_{(\cdot)} : C^k \to \mathbb{Q} \cap [0,1]$$

A base element is non-empty, iff $\forall \gamma \in C^k\ x_\gamma < y_\gamma$; and then it will contain the function $f_0 : C^\omega \to [0,1]$ defined via $f_0(p) = \frac{1}{2}\left(x_{p_{\leq k}} + y_{p_{\leq k}}\right)$, which clearly depends only on the prefix of length k of its argument.

3. As the actions of the players beyond the finite prefix determining the outputs is irrelevant, and taking into consideration the definition of the subgame-perfect preferences, the claim is that any finite game in extensive form has a subgame perfect equilibrium. This well-known result (KUHN [16]) is easily proven by backwards induction: let the players who move last pick an optimal (for them) choice. Then the players who move second-but-last have guaranteed outcomes associated with their moves, so they can optimize, and so on.

□

Corollary 10 ([6])**.** Any sequential game with continuous payoff functions and finitely many choices has a subgame perfect Nash equilibrium.

We shall consider a number of variations/extensions. First, we investigate stochastic[7] infinite sequential games with continuous payoff functions as a variation of Definition 4: a stochastic infinite sequential game with real-valued payoffs is a tuple $\langle A, C, d, (f_a)_{a\in\mathbb{N}}, P\rangle$ where A, C, f_a are as above, $\mathbf{n} \notin A$, $d : C^* \to A \cup \{\mathbf{n}\}$ and P assigns a probability distribution over C to each $w \in d^{-1}(\{\mathbf{n}\})$.

The notion of the play induced by a strategy profile is replaced by a probability distribution over plays induced by a strategy profile: essentially, we choose according to the strategy profile in vertices controlled by a player, and stochastically according to $P(w)$ in vertices controlled by $\mathbf{n}$ (nature). A player prefers a strategy profile to another one if the expected value of his payoff function regarding the former induced probability distribution exceeds the expected value regarding the latter. We obtain a notion of a subgame perfect equilibrium as before.

Lemma 11. Consider stochastic infinite sequential games with continuous payoff-functions and finite choices sets. We find:

1. The subgame-perfect preferences produce a compact-strategies, open-preference class $\mathcal{SS}$.
2. The games with payoffs fully determined after finitely many moves are a dense subset $\mathcal{SS}^f \subseteq \mathcal{SS}$.
3. All games in $\mathcal{SS}^f$ have a Nash equilibrium.

Proof. 1. For a synthetic approach to continuity on probability measures etc, see [6]. In particular, the map from strategy profiles to induced probability distributions is easily seen to be continuous; and integration is as well. The remaining argument proceeds as in Lemma 9 (1).

2. The stochastic vertices do not impact the argument, so it works as in Lemma 9 (2).

3. As in Lemma 9 (3), we use backwards inductions: the value of a given leaf for a player is immediately obtained from the payoff function. In any vertex controlled by a player, he picks a choice guaranteeing him the optimal value, and the value of that vertex for any

[6]This extends [10, Corollary 4.2] from finitely many players to countably many players.

[7]In this, we are answering a question raised by ASSALÉ ADJÉ at CSL-LiCS.

player p is identical to the value of the corresponding child for p. The value of a nature vertex is the expected value of its children according to the probability distribution given by P. This is easily seen to yield a subgame perfect equilibrium.

□

Corollary 12. Stochastic infinite sequential games with continuous payoff-functions and finite choices sets have subgame-perfect equilibria.

In a similar fashion we could consider multi-player multi-outcome Blackwell games [1, 25] with continuous payoff functions (into $[0,1]$). In a Blackwell game, the players act jointly and probabilistic in each round and thus determine the next vertex in the tree. In particular, we are now dealing with a tree over C^A instead of just C; i.e. d is no longer relevant, strategies are functions $s_a : (C^A)^* \to P(C)$, where $P(C)$ shall denote the set of probability distributions over C, strategy profiles are of the form $s : A \times (C^A)^* \to P(C)$ and payoff functions go $f_a : (C^A)^\omega \to [0,1]$. Each strategy profile induces a probability distribution over $(C^A)^\omega$ (also from a given vertex onwards), payoffs for players are then again the expected values of their payoff functions given the induced probability distribution.

Lemma 13. Consider Blackwell games with continuous payoff-functions, countably many agents and a finite choice set. We find:

1. The subgame-perfect preferences produce a compact-strategies, open-preference class $\mathcal{B}$.
2. The games with payoffs fully determined after finitely many moves and depending only on the choices made by finitely many players are a dense subset $\mathcal{B}^f \subseteq \mathcal{B}$.
3. All games in $\mathcal{B}^f$ have a Nash equilibrium.

Proof.

1. Very much like Lemma 11 (1). The induced probability distribution depends continuously on the strategy profile, and the expected value is a continuous map. By Tychonoffs theorem it is unproblematic that C is replaced by C^A now, as this space is still compact.
2. As in Lemma 9 (2), we only need to argue that the continuous payoff functions determined by their values on $(C^{A_0})^n$ for finite $A_0 \subseteq A$ and $n \in \mathbb{N}$ are dense in $\mathcal{C}((C^A)^\omega, [0,1])$. This in turn follows just as before, taking into account the definition of the product topology.
3. We can ignore the players whose actions do not impact the payoff functions, and the moves in the game taking place after the payoffs have been determined. Thus we are dealing with a finite tree, and can use backwards induction again. The leaves of the finite tree are assigned payoff values from the payoff functions. The preceding layer of vertices now corresponds to usual normal form games (with finitely many players and finitely many choices), who have Nash equilibria by Nash' classic result [27]. Pick a Nash equilibrium for each such vertex, and assign it the corresponding outcomes as values. Then the layer preceding this one can be handled in the same way, etc. – so the claim follows by backwards induction.

□

Corollary 14 ([8])**.** Any multi-outcome Blackwell game with countably many players and continuous payoff functions has a subgame-perfect Nash equilibrium.

Concurrent games may be another interesting class of games to apply this approach to. Determinacy questions of these were discussed in [13, 14].

[8] This extends [10, Theorem 6.1] from finitely many players to countably many players.

4 From discrete to real-valued payoffs

A rather simple argument allows us to transfer existence theorems for equilibria in games with finitely many players, from Borel-measurable valuations over finitely many outcomes to Borel-measurable real-valued payoff functions, if one is willing to replace the original notions by their ε-counterparts. If $v : \mathbf{S} \to [0,1]^m$ is the Borel-measurable payoff function, then for any $k \in \mathbb{N}$ we define $v_{(k)} : \mathbf{S} \to \{1, \dots, (k+1)\}^m$ by $v_{(k)}^{-1}(i_1, \dots, i_m) := v^{-1}\left([\frac{i_1-1}{k}, \frac{i_1}{k}) \times \dots \times [\frac{i_m-1}{k}, \frac{i_m}{k})\right)$. Then any $v_{(k)}$ is again a Borel measurable valuation. Furthermore, we define the preferences $\prec_n$ for the n-th player by $(i_1, \dots, i_m) \prec_n (j_1, \dots, j_m)$ iff $i_n < j_n$. Now any Nash equilibrium of the resulting game is a $\frac{1}{k}$-Nash equilibrium of the original game, and every subgame perfect equilibrium of the resulting game is a $\frac{1}{k}$-subgame perfect equilibrium of the original game. Note that the same argument can be used for games with countably many players and countably many outcomes, without any restriction on the branching of the game-tree.

Corollary 15. [9] Sequential games with countably many players and Borel-measurable payoff functions with upper-bounds admit ε-Nash equilibria.

Proof. By combining the statement of [19, Theorem 3.2] with the argument above. □

5 Infinite sequential games with infinitely ascending preferences

As soon as the continuity requirement for the payoff function (or, more generally, the openness of the preferences) is dropped, Nash equilibria may fail to exist. We provide a generic folklore counterexample, and will proceed to demonstrate that the underlying feature is essential for the failure of existence of Nash equilibria. The counterexample only requires a single player, and its payoff function is in a sense *the least discontinuous* payoff function, and in particular is Δ_2^0-measurable.

Example 16. Let the payoff function $P : \{0,1\}^{\mathbb{N}} \to [0,1]$ for the single player be defined by $P(1^n 0 p) = \frac{n}{n+1}$ for all $p \in \{0,1\}^{\mathbb{N}}$ and $P(1^{\mathbb{N}}) = 0$. As P does not attain its supremum, the resulting game cannot have a Nash equilibrium.

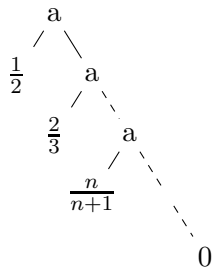

We proceed to show in particular that the presence of a converging sequence of plays $(p^n)_{n\in\mathbb{N}}$ such that a player prefers p^{n+1} to p^n, but prefers any p^n to $\lim_{i\to\infty} p^i$, is a crucial feature of the example above to have no Nash equilibrium. The proof will be an adaption of the main result of [19] by the first author. Under the additional assumption of antagonistic preferences in a two-player game, we can even obtain subgame perfect equilibria.

[9] In his survey [26], MERTENS sketches an observation by himself and NEYMAN that one may use Borel determinacy to directly obtain this result.

In this section the preferences of the players are restricted to strict weak orders, so we recall their definition below.

Definition 17 (Strict weak order)**.** A relation $\prec$ is called a *strict weak order* if it satisfies:

$$\begin{array}{ll}\forall x, & \neg(x \prec x)\\ \forall x,y,z, & x \prec y \wedge y \prec z \Rightarrow x \prec z\\ \forall x,y,z, & \neg(x \prec y) \wedge \neg(y \prec z) \Rightarrow \neg(x \prec z)\end{array}$$

Definition 35 below slightly rephrases Definitions 2.3 and 2.5 from [19]: the guarantee of a player is the smallest set of outcomes that is upper-closed w.r.t. the strict-weak-order preference of the player and includes every incomparability classe (of the preference) that contains any outcome compatible with a given strategy of the player in the subgame at a given node of a given infinite sequential game. The best guarantee of a player consists of the intersection of all her guarantees over the set of strategies.

Definition 18 (Agent (best) guarantee)**.** Let $\langle A, C, d, O, v, (\prec_a)_{a\in A}\rangle$ be a game where the $\prec_a$ are strict weak orders.

$$\begin{array}{l}\forall a \in A, \forall \gamma \in C^*, \forall s : d^{-1}(a) \to C, \quad g_a(\gamma, s) :=\\ \{o \in O \mid \exists p \in P(s|_{\gamma C^\omega}) \cap \gamma C^\omega, \neg(o \prec_a v(p))\}\\ G_a(\gamma) := \bigcap_s g_a(\gamma, s)\end{array}$$

We write $g_a(s)$ and G_a instead of $g_a(\gamma, s)$ and $G_a(\gamma)$ when γ is the empty word.

Lemma 2.4. from [19] still holds without major changes in the proofs, so we do not display it, but note that when speaking about $\prec_a$-terminal intervals (which are upper-closed sets), we now actually refer to the terminal intervals of the lift of $\prec_a$ from outcomes to the equivalence classes of outcomes induced by the strict weak order. Also, we collect some more useful facts in Observation 19 below.

Observation 19. Let $\langle A, C, O, d, v, (\prec_a)_{a\in A}\rangle$, let $a \in A$, assume that $\prec_a$ is a strict weak order, and let $\gamma \in C^*$.

1. $d(\gamma) \neq a \Rightarrow G_a(\gamma) = \cup_{c\in C} G_a(\gamma \cdot c)$
2. $d(\gamma) = a \Rightarrow G_a(\gamma) = \cap_{c\in C} G_a(\gamma \cdot c)$
3. $d(\gamma) = a \wedge |C| < \infty \Rightarrow \exists c \in C,\ G_a(\gamma) = G_a(\gamma \cdot c)$

Proof. For example, for 2. note that $G_a(\gamma) = \cap_s g_a(\gamma, s) = \cap_{c\in C} \cap_{s(\gamma)=c} g_a(\gamma, s) = \cap_{c\in C} \cap_s g_a(\gamma \cdot c, s)$
$= \cap_{c\in C} G_a(\gamma \cdot c)$. □

This section's proofs of existence of equilibrium rely on each player having a (minimax-style) optimal strategy if all other players team up against her. Lemma 20 below provides a sufficient condition for such strategies to exist, *i.e.* for the best guarantee to be witnessed.

Lemma 20. Let $\langle A, C, O, d, v, (\prec_a)_{a\in A}\rangle$ be a game where C is finite, let $a \in A$, and let us assume the following.

1. $\prec_a$ is a strict weak order.
2. For every play $p \in C^\omega$, increasing $\varphi : \mathbb{N} \to \mathbb{N}$, and sequence $(s_n)_{n\in\mathbb{N}}$ of strategies for Player a, if $g_a(p_{<\varphi(n+1)}, s_{n+1}) \subsetneq g_a(p_{<\varphi(n)}, s_n)$ for all $n \in \mathbb{N}$, then $v(p) \in \cap_{n\in\mathbb{N}} g_a(p_{<\varphi(n)}, s_n)$.

Then for all $\gamma \in C^*$ there exists $s \in S_a$ such that $g_a(\gamma, s) = G_a(\gamma)$.

Proof. Wlog we only prove that there exists $s \in S_a$ such that $g_a(s) = G_a$, *i.e.* where the γ from the claim is the empty word. Let $s_0 : d^{-1}(a) \to C$ be a strategy for Player a and let us build inductively a sequence $(s_n)_{n\in\mathbb{N}}$ of strategies for Player a, as follows, where case 3. implicitly invokes Observation 19.

- Let $s_{n+1}|_{C^{<n}} := s_n|_{C^{<n}}$.
- For all $\gamma \in C^n \backslash d^{-1}(a)$, let $s_{n+1}|_{\gamma C^*} := s_n|_{\gamma C^*}$.
- For all $\gamma \in C^n \cap d^{-1}(a)$,
  1. if $g_a(\gamma, s_n) \subseteq G_a$ then let $s_{n+1}|_{\gamma C^*} := s_n|_{\gamma C^*}$,
  2. if $G_a \subsetneq g_a(\gamma, s_n)$ and there exists $\mu : d^{-1}(a) \cap \gamma C^* \to C$ such that $g_a(\gamma, \mu) \subseteq G_a$, let $s_{n+1}|_{\gamma C^*} := \mu$,
  3. otherwise[10] let $s_{n+1}(\gamma) := c$ such that $G_a(\gamma{\cdot}c) = G_a(\gamma)$, and let $s_{n+1}|_{\gamma CC^*} := s_n|_{\gamma CC^*}$.

Let s be the limit strategy of the sequence $(s_n)_{n\in\mathbb{N}}$ and first note that, using Observation 19, one can prove by induction on γ that $G_a(\gamma) \subseteq G_a$ for every $\gamma \in C^*$ that is compatible with s. Next, let $p \in P(s)$ be a path compatible with s. If p has a prefix γ that fell into Cases 1. or 2. during the recursive construction above, then $v(p) \in G_a$, so let us now assume that case 3. applies at every node $p_{<n} \in d^{-1}(a)$. If such nodes are finitely many, let $p_{<n}$ be the deepest one, so $d(p_{n+1+k}) \neq a$ for all $k \in \mathbb{N}$, and $v(p) \in g_a(p_{n+1}, t)$ for all strategies t for a, so $v(p) \in G_a(p_{<n+1}) = \bigcap_t g_a(p_{<n+1}, t)$. So $v(p) \in G_a$ since $G_a = G_a(p_{<n+1})$ by Case 3. Let us now assume that such nodes are infinitely many. If $G_a(p_{<n}) \subsetneq G_a$ for some $p_{<n} \in d^{-1}(a)$, there exists $\mu : d^{-1}(a) \cap p_{<n}C^* \to C$ such that $G_a(p_{<n}) \subseteq g_a(p_{<n}, \mu) \subsetneq G_a$ since $G_a(p_{<n}) = \bigcap_t g_a(p_{<n}, t)$ by definition, which would mean that Case 1. or 2. applies; so $G_a(p_{<n}) = G_a$ for all $p_{<n} \in d^{-1}(a)$, and subsequently for all n. Also, the best guarantee is never witnessed (through Case 2.) at any node $p_{<n} \in d^{-1}(a)$, and subsequently for all n. If $v(p) \notin G_a$, the previous two remarks allow us to build inductively a sequence $(t_n)_{n\in\mathbb{N}}$ of strategies for a such that $v(p) \notin g_a(p_{<0}, t_0)$ and $g_a(p_{<n+1}, t_{n+1}) \subsetneq g_a(p_{<n}, t_n)$ for all $n \in \mathbb{N}$, which would imply $v(p) \in \cap_{n\in\mathbb{N}} g_a(p_{<n}, t_n)$ by assumption of the lemma, contradiction. □

More specifically, this section's proofs of existence of equilibrium rely on each player having a strategy that is (minimax-style) optimal for every subgame at once. Lemma 21 below gives a sufficient condition for such strategies to exists. It relies on Condition 3, which is also the conclusion in Lemma 20, to perform a key quantifier inversion. Note that Condition 2 in Lemma 21 is weaker than that in Lemma 20, and that finiteness of C is used in Lemma 20 only. It is the main reason why Lemmas 20 and 21 are not merged.

Lemma 21. Let $\langle A, C, O, d, v, (\prec_a)_{a\in A}\rangle$ be a game, let $a \in A$, and let us assume the following.

1. $\prec_a$ is a strict weak order.
2. For every play $p \in C^\omega$ and increasing $\varphi : \mathbb{N} \to \mathbb{N}$, if $d(p_{<\varphi(n)}) = a$ and $G_a(p_{<\varphi(n+1)}) \subsetneq G_a(p_{<\varphi(n)})$ for all $n \in \mathbb{N}$, then $v(p) \in \cap_{n\in\mathbb{N}} G_a(p_{<\varphi(n)})$.
3. For all $\gamma \in C^*$ there exists $s \in S_a$ such that $g_a(\gamma, s) = G_a(\gamma)$.

Then there exists s such that $g_a(\gamma, s) = G_a(\gamma)$ for all $\gamma \in C^*$.

[10] Note that due to the properties of a strict weak order, the sets of the form $g_a(\gamma, s)$ and $G_a(\gamma)$ are linearly ordered by inclusion $\subseteq$. Thus, $G_a \subsetneq g_a(\gamma, s_n)$ holds in this case, too.

Proof. We proceed similarly as in the proof of Lemma 20. Let s_0 be a strategy for Player a and let us build inductively a sequence $(s_n)_{n\in\mathbb{N}}$ of strategies for Player a. The recursive definition below is different from the one in the proof of Lemma 20 in three respects: the three occurrences of G_a in Cases 1. and 2. are replaced with $G_a(\gamma)$. Case 3. is deleted since it never applies by assumption. Finally, two inclusions are replaced with equalities.

- Let $s_{n+1}|_{C^{<n}} := s_n|_{C^{<n}}$
- For all $\gamma \in C^n \backslash d^{-1}(a)$, let $s_{n+1}|_{\gamma C^*} := s_n|_{\gamma C^*}$.
- For all $\gamma \in C^n \cap d^{-1}(a)$,
  1. if $g_a(\gamma, s_n) = G_a(\gamma)$ then let $s_{n+1}|_{\gamma C^*} := s_n|_{\gamma C^*}$,
  2. if $G_a(\gamma) \subsetneq g_a(\gamma, s_n)$, let $s_{n+1}|_{\gamma C^*} := \mu$ where $\mu : d^{-1}(a) \cap \gamma C^* \to C$ is such that $g_a(\gamma, \mu) = G_a(\gamma)$.

Let s be the limit strategy of the sequence $(s_n)_{n\in\mathbb{N}}$ and first note that, using Observation 19, one can prove by induction on γ that $G_a(\gamma) \subseteq G_a$ for every $\gamma \in C^*$ that is compatible with s. Next, let $p \in P(s)$ be a path compatible with s. Due to the uniformity of the recursive definition, it suffices to show that $v(p) \in G_a$ to prove the full statement.

If Case 2. applies only finitely many times in the construction of s, the sequence $(s_n|_{\{\gamma\in C^* \mid p\in\gamma C^\omega\}})_{n\in\mathbb{N}}$ is eventually constant, so $v(p) \in g_a(p_{<n}, s_n) = G_a(p_{<n}) \subseteq G_a$ for some n. Otherwise, there exists an increasing function $\varphi : \mathbb{N} \to \mathbb{N}$ with $d(p_{<\varphi(n)}) = a$ and $G_a(p_{<\varphi(n+1)}) \subsetneq G_a(p_{<\varphi(n)})$ for all $n \in \mathbb{N}$, so $v(p) \in \cap_{n\in\mathbb{N}} G_a(p_{<\varphi(n)}) \subseteq G_a(p_{<\varphi(0)}) \subseteq G_a$. □

Theorems 23 and 24 below both prove the existence of subgame perfect equilibria for antagonistic games, either when the choice set C is finite, or when the outcome set O is finite. Since their proofs are similar, most is factored out in Lemma 22 below.

Lemma 22. Let $\langle\{a, b\}, C, O, d, v, \{\prec, \prec^{-1}\}\rangle$ be a two-player game. Let $\Gamma \subseteq \mathcal{P}(C^\omega)$ and assume the following.

1. $\prec$ is a strict weak order.
2. For every play $p \in C^\omega$ and increasing sequence $\varphi : \mathbb{N} \to \mathbb{N}$, if $d(p_{<\varphi(n)}) = a$ and $G_a(p_{<\varphi(n+1)}) \subsetneq G_a(p_{<\varphi(n)})$ for all $n \in \mathbb{N}$, then $v(p) \in \cap_{n\in\mathbb{N}} G_a(p_{<\varphi(n)})$.
3. For all $\gamma \in C^\omega$, there exists s such that $g_a(\gamma, s) = G_a(\gamma)$ (resp. $g_b(\gamma, s) = G_b(\gamma)$).
4. For all non-empty closed $E \subseteq C^\omega$, there are $\prec$-extremal elements in $v[E]$.
5. For every $\prec$-extremal interval I and $\gamma \in C^*$, we have $\left(v^{-1}[I] \cap \gamma C^\omega\right) \in \Gamma$.
6. The game $\langle C, D, W\rangle$ is determined for all $W \in \Gamma$, $D \subseteq C^*$.

Then the game $\langle\{a, b\}, C, O, d, v, \{\prec, \prec^{-1}\}\rangle$ has a subgame perfect equilibrium.

Proof. By invoking Lemma 21 once for Player a and once for Player b, let us build a strategy profile $s : C^* \to C$, such that $g_X(\gamma, s_X) = G_X(\gamma)$ for all $\gamma \in C^*$ and $X \in \{a, b\}$. Let $\gamma \in C^*$ and let us prove that $G_a(\gamma) \cap G_b(\gamma) = \{\min_<(G_a(\gamma))\} = \{\max_<(G_b(\gamma))\}$. Consider the game $\langle C, D, W\rangle$ (as in Definition 5) where the winning set is defined by $W := \{\alpha \in \gamma C^\omega \mid v(\alpha) \in G_a(\gamma)\backslash\{\min_<(G_a(\gamma))\}\}$ and where Player a owns exactly the nodes in $D := (C^* \setminus \gamma C^*) \cup (d^{-1}(\{a\}) \cap \gamma C^*)$. By Assumption 5 the set W is in Γ, so by Assumption 6 the game $\langle C, D, W\rangle$ is determined. By definition of the best guarantee, Player a has no winning strategy for this game, so Player b has a winning strategy, which means that $G_b(\gamma) \subseteq \{\min_<(G_a(\gamma))\} \cup O\backslash G_a(\gamma)$.

Since $G_a(\gamma) \cap G_b(\gamma)$ must be non-empty, otherwise the two guarantees are contradictory, $G_a(\gamma) \cap G_b(\gamma) = \{\min_<(G_a(\gamma))\}$. This means that the subprofile of s rooted at γ induces the outcome $\min_<(G_a(\gamma))$ (which equals $\max_<(G_b(\gamma))$ by symmetry), and it is optimal for both players. □

Theorem 23. Let $\langle \{a,b\}, C, O, d, v, \{\prec, \prec^{-1}\}\rangle$ be a two-player antagonistic game, where C is finite. Let $\Gamma \subseteq \mathcal{P}(C^\omega)$ and assume the following.

1. $\prec$ is a strict weak order.
2. For every $p \in C^\omega$, sequence $(s_n)_{n\in\mathbb{N}}$ of strategies for $X \in \{a,b\}$, and increasing $\varphi : \mathbb{N} \to \mathbb{N}$, if $d(p_{<\varphi(n)}) = a$ and $g_X(p_{<\varphi(n+1)}, s_{n+1}) \subsetneq g_X(p_{<\varphi(n)}, s_n)$ for all $n \in \mathbb{N}$, then $v(p) \in \cap_{n\in\mathbb{N}} g_X(p_{<\varphi(n)}, s_n)$.
3. For every $\prec$-extremal interval I and $\gamma \in C^*$, we have $\left(v^{-1}[I] \cap \gamma C^\omega\right) \in \Gamma$.
4. The game $\langle C, D, W\rangle$ is determined for all $W \in \Gamma$, $D \subseteq C^*$.

Then the game $\langle \{a,b\}, C, O, d, v, \{\prec, \prec^{-1}\}\rangle$ has a subgame perfect equilibrium.

Proof. By application of Lemma 22. (Note that Condition 2 of Theorem 23 is an "upper bound" of Condition 2 of Lemma 20 and Condition 2 of Lemma 21.) Condition 3 is proved by Lemma 20. For Condition 4, let E be a non-empty closed subset of C^ω, and let T be the tree such that $[T] = E$. Consider the game where Player a plays alone on T. Since Player a can maximise her best guarantee by Lemma 20, and since all her guarantees are singletons, $v[E]$ has a $\prec$-maximum. Likewise, it has a $\prec$-minimum, by considering Player b. □

Theorem 24. Let $\langle \{a,b\}, C, d, O, v, \{<, <^{-1}\}\rangle$ be an infinite sequential game where O is finite and $<$ is a strict linear order. Let $\Gamma \subseteq \mathcal{P}(C^\omega)$ and assume the following.

1. $\forall O' \subseteq O, \forall \gamma \in C^*, \quad \{\alpha \in C^\omega \mid v(\gamma\alpha) \in O'\} \in \Gamma$
2. The game $\langle C, D, W\rangle$ is determined for all $W \in \Gamma$ and $D \subseteq C^*$.

Then the game $\langle \{a,b\}, C, d, O, v, \{<, <^{-1}\}\rangle$ has a subgame perfect equililbrium.

Proof. by Lemma 22 where Conditions 2, 3, and 4 hold by finiteness of O. □

Corollary 25. Let $\langle \{a,b\}, C, d, O, v, \{<, <^{-1}\}\rangle$ be an infinite sequential game where O is finite and $<$ is a strict linear order. If $v^{-1}(o)$ is quasi-Borel for all $o \in O$, the game has a subgame perfect equilibrium.

Proof. From Theorem 24, quasi-Borel determinacy [25], and Lemma 3.1. in [19]. □

Corollary 26. A sequential two-player zero-sum game with Borel measurable payoffs has ε-subgame perfect equilibria.

Proof. By Corollary 25 and the argument in Section 4. □

In the remainder of this section we discuss existence of Nash equilibria in multi-player games. Theorem 27 below is our most general result.

Theorem 27. Let $\langle A, C, O, d, v, (\prec_a)_{a\in A}\rangle$ be a game, let $\Gamma \subseteq \mathcal{P}(C^\omega)$, and assume the following.

1. The $\prec_a$ are strict weak orders.
2. The game $\langle C, D, W\rangle$ is determined for all $W \in \Gamma$, $D \subseteq C^*$.

3. For every $a \in A$ and $\prec_a$-terminal interval I and $\gamma \in C^*$, we have $\left(v^{-1}[I] \cap \gamma C^\omega\right) \in \Gamma$.
4. For every play $p \in C^\omega$ and increasing sequence $\varphi : \mathbb{N} \to \mathbb{N}$, if $d(p_{<\varphi(n)}) = a$ and $G_a(p_{<\varphi(n+1)}) \subsetneq G_a(p_{<\varphi(n)})$ for all $n \in \mathbb{N}$, then $v(p) \in \cap_{n\in\mathbb{N}} G_a(p_{<\varphi(n)})$.
5. For all $\gamma \in C^\omega$, there exists s such that $g_a(\gamma, s) = G_a(\gamma)$.

Then the game $\langle A, C, O, d, v, (\prec_a)_{a\in A}\rangle$ has a Nash equilibrium.

Proof. Since the proof is similar to that of Theorem 2.9 in [19], we rephrase and give it a more intuitive flavour. Let σ be a strategy profile where every player is using a witness to Lemma 21. Let p be the induced play. We now turn σ into a Nash equilibrium with p as induced play by use of threats. More specifically, at each node $p_{<n}$ we let the players other than $a := d(p_{<n})$ threaten Player a that if she deviates from p exactly at $p_{<n}$, they will team up against her at every subsequent position γ after $p_{<n}$ other than those extending the prescribed $p_{<n+1}$.

We claim that if they team up, they can prevent Player a from getting better outcome than $v(p)$ by deviating to γ, which will suffice. Let us build a win-lose game $\langle C, D, W\rangle$, with Player a against her threatening opponents gathered as a meta-player, and where the winning set for Player a is defined by $W = v^{-1}[I] \cap \gamma C^\omega$, where $I := \{o \in O \mid v(p) \prec_a o\}$, and D is defined by $D = d^{-1}(\{a\}) \cup (C^* \setminus \gamma C^*)$. This game is determined by Assumptions 2 and 3, and Player a looses it, otherwise her winning strategy would guarantee that $v(p) \notin G_a(p_{<n})$ and thus contradict the choice of p. Therefore the threat of the opponents of Player a is effective. □

Theorem 28 below is a simpler version of Theorem 27 that does only involve primitive notions from the definition of a game. Especially, it does not refer to the notion of guarantee. Via a necessary and sufficient condition, it shows how essential the feature of Example 16 is for the existence of Nash equilibrium.

Theorem 28. Let g be a $\langle A, C, O, d, v, (\prec_a)_{a\in A}\rangle$ be a game where the $\prec_a$ are strict weak orders and v is Borel-measurable. The following are equivalent.

1. For every $X \in \{a, b\}$ and $(p^n)_{n\in\mathbb{N}}$ sequence of plays in C^ω converging to some p, and increasing $\varphi : \mathbb{N} \to \mathbb{N}$, if for all $n \in \mathbb{N}$ we have $d(p_{<\varphi(n)}) = X$, $p_{<\varphi(n)} = p^n_{<\varphi(n)}$, $p_{\varphi(n)} \neq p^n_{\varphi(n)}$, and $v(p^n) \prec_X v(p^{n+1})$, then $v(p^n) \prec_X v(p)$ for all $n \in \mathbb{N}$.
2. Every finite-branching game derived from the original game by pruning has an NE.

Proof. Let us first prove 1. ⇒ 2. by invoking Theorem 27. More specifically, let T be a finite-branching, infinite subtree of C^ω and consider the restriction of the original game to T. Conditions 3 and 2 follow from Borel measurability and [24]. Condition 5 comes from Lemma 20 (actually a straightforward extension of Lemma 20 to trees with finite-yet-unbounded branching), and Condition 4 follows directly from the assumption.

For 2. ⇒ 1., let $X \in \{a, b\}$ and let $(p^n)_{n\in\mathbb{N}} \to p \in C^\omega$ and increasing $\varphi : \mathbb{N} \to \mathbb{N}$ such that for all $n \in \mathbb{N}$ we have $d(p_{<\varphi(n)}) = X$, $p_{<\varphi(n)} = p^n_{<\varphi(n)}$, $p_{\varphi(n)} \neq p^n_{\varphi(n)}$, and $v(p^n) \prec_X v(p^{n+1})$. Let T be the tree made of the prefixes of p and the p^n. Since the game induced by T has an NE and its tree structure is similar to Example 16, $v(p^n) \prec_X v(p)$ must hold for all $n \in \mathbb{N}$. □

However, the modification of Example 16 below shows that the conditions of Theorem 28 are not necessary for the mere existence of Nash equilibria.

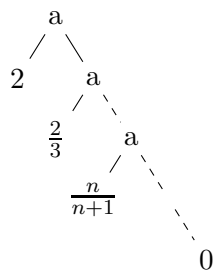

For further comparison, the preparatory work before [19, Theorem 2.9] considers strict well-orders only; then [19, Theorem 2.9] considers strict well-founded orders, since linear extensions of these make it possible to invoke the special, linear case, knowing that any Nash equilibrium for these extensions is still a Nash equilibrium for the original preferences. However, let us explain why Theorem 27 assumes that preferences are strict weak orders, instead of more general strict partial orders. In the preparatory work before both results, the algorithm that builds a play step by step makes decisions based on the guarantees that the subgames offer. If the guarantees of one player were not ordered by a strict weak order, the player might eventually regret a previous decision, in the same way that backward induction on partially ordered preferences may not yield a Nash equilibrium (see *e.g.*, [15] for a concrete example or page 3 of [18] for a generic one). So the algorithm has to run on strict weak orders. (In [19, Theorem 2.9] it even runs on strict linear orders.)

If we wanted to consider strict partial orders and extend them linearly for the algorithm to work, we would potentially run into two problems: first, there may not exist any linear extension preserving Condition 4. Second, assumptions 2 and 3 of Theorem 27 make sure that the win-lose games associated with the $\prec_a$-terminal intervals are determined, which is a requirement for the proof to work. If the preferences were not strict weak orders, we might think of replacing the condition on terminal intervals by a condition on the upper-closed sets and then extend the preferences linearly for the algorithm to work, but in the special case where the preference of one player were the empty relation, every subset would be an upper-closed set and its preimage by v would be in the pointclass with nice closure property, by assumption. If, in addition, each outcome is mapped to at most one play, it implies that each subset of C^ω is in the pointclass, so Theorem 27 could be used with the axiom of determinacy only, but not with, *e.g.*, Borel determinacy. On the contrary, [19, Theorem 2.9, Assumption 3] is not an issue since there are only countably many outcomes in that setting.

Theorem 27 has a corollary pertaining to sequential games with real-valued payoffs. Rather than the usual Euclidean topology, we consider the lower topology generated by $\{(-\infty, a) \mid a \in \mathbb{Q}\}$. This space will be denoted by $\mathbb{R}_>$. Note that continuous functions with codomain $\mathbb{R}_>$ are often called upper semi-continuous. As $\mathrm{id} : \mathbb{R}_> \to \mathbb{R}$ is complete for the Σ^0_2-measurable functions [42, 40], we see that the Borel sets[11] on $\mathbb{R}_>$ are the same as the Borel sets on $\mathbb{R}$. Moreover, if $(p^n)_{n\in\mathbb{N}}$ is a converging sequence of plays, and $P : \{0,1\}^{\mathbb{N}} \to \mathbb{R}_>$ is a continuous payoff function, then $P(\lim_{i\in\mathbb{N}} p^i) \geq \limsup_{i\in\mathbb{N}} P(p^i)$. In particular, Condition 4 in Theorem 27 is always satisfied for the preferences obtained from upper semi-continuous payoff functions.

Corollary 29. Sequential games with countably many players, finitely many choices and upper semi-continuous payoff functions have Nash equilibria.

[11] As $\mathbb{R}_>$ is not metric (but still countably based), the definition of the Borel hierarchy has to be modified as demonstrated by Selivanov [37]. A move towards definitions of Borel measurability on even more general spaces can be found in [33].

6 On the existence of Pareto-optimal NE

In this section we investigate very general classes of games that guarantee existence of Nash equilibria, and such that there exists an NE that is Pareto-optimal among all the profiles of the game (not just Pareto-optimal among all Nash equilibria). In the following, we shall assume that any outcome is realizable to avoid unnecessary case-distinctions.

Lemma 30. Let Γ be a determined pointclass. Then every infinite sequential two-player game with a Γ-measurable outcome function and preferences $y \prec_a x_1 \prec_a \ldots \prec_a x_n$ and $y \prec_b x_n \prec_b \ldots \prec_b x_1$, has a Pareto-optimal NE.

Proof. By assumption that every outcome is realizable, there is some path p through the game yielding a payoff that is not y. For each vertex along this path, by determinacy either the opponent can enforce the outcome y, or the controller can enforce some upper interval. As long as the opponent can enforce y, he can force the controller to play along the chosen path by threaten punishment by y for deviation. If we ever reach a vertex where the controller (w.l.o.g. a) can enforce $\{x_1, \ldots, x_n\}$, there will be some minimal upper set $\{x_i, \ldots, x_n\}$ (from her perspective) that she can enforce. By determinacy, again, the opponent can enforce $\{y, x_1, \ldots, x_i\}$. We then let both players play their enforcing strategy from this node onwards.

The constructed partial strategies can be extended in an arbitrary way to yield a Nash equilibrium with another outcome than y, and these are all Pareto-optimal. □

Theorem 31. We fix a non-empty set of players A and a non-empty set of outcomes O. Let Γ be a determined pointclass closed under rescaling and union with clopens. Then the following are equivalent for a family $(\prec_a)_{a\in A}$ of linear preferences:

1. The inverse of the preferences are well-founded and $\forall a, b \in A, \forall x, y, z \in O, \neg(z \prec_a y \prec_a x \wedge x \prec_b z \prec_b y)$.
2. Every finite sequential game (built from A, O, $(\prec_a)_{a\in A}$) with three leaves has a Pareto-optimal NE.
3. Every infinite sequential game (built from A, O, $(\prec_a)_{a\in A}$) with a Γ-measurable outcome function has a Pareto-optimal NE.

Proof. 3. $\Rightarrow$ 2. Clear.

2. $\Rightarrow$ 1. By contraposition, let us assume that $z \prec_a y \prec_a x$ and $x \prec_b z \prec_b y$, and note that the game below has only one NE yielding outcome z.

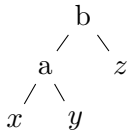

1. $\Rightarrow$ 3. By [20, Lemma 25] the second assumption in 1. implies that there exists a partition $\{O_i\}_{i\in I}$ of O and a linear order $<$ over I such that $i < j$ implies $x <_a y$ for all $a \in A$ and $x \in O_i$ and $y \in O_j$, and such that $<_b|_{O_i} = <_a|_{O_i}$ or $<_b|_{O_i} = <_a|_{O_i}^{-1}$ for all $a, b \in A$. By the well-foundedness assumption, I has a $<$-maximum m.

   Fix some $a \in A$, let $\{x_1, \ldots, x_n\} := O_m$ (again, by well-foundedness, each slice is finite) such that $x_n <_a \cdots <_a x_1$, let $A_0 := \{b \in A | <_b |_{O_m} = <_a |_{O_m}\}$, let $A_1 := A \backslash A_0$, and

let $y \notin O_m$. Let us derive a new game on the same tree: each vertex of the original game owned by $b \in A$ is now owned by A_0 if $b \in A_0$ and by A_1 otherwise. Each play of the original game that induces an outcome outside of O_m induces y in the derived game. The new preferences are $y <_{A_0} x_n <_{A_0} \cdots <_{A_0} x_1$ and $y <_{A_1} x_1 <_{A_1} \cdots <_{A_1} x_n$. By Lemma 30, the derived has a Pareto-optimal NE (which cannot yield y, as this is the only non-Pareto-optimal outcome). It is also a Pareto-optimal NE for the original game.

□

The situation for non-linear orders is less clear. Certainly, whenever some linearization avoids the forbidden pattern from Theorem 31 (1.), there will be a Pareto-optimal NE (as being Pareto-optimal w.r.t. the linearization implies being Pareto-optimal w.r.t. the original preferences). However, we do not know whether partial preferences such that any linearization has the forbidden pattern is enough to enable absence of Pareto-optimal NE. Two examples that could potentially play a similar role to the generic counterexample in Theorem 31 (2. $\to$ 1.) follow:

Example 32. We consider a finite two-player game with outcomes $\{x, y, z, \alpha, \beta, \gamma\}$, preferences $\gamma \prec_a y \prec_a x$ and $z \prec_a \beta \prec_a \alpha$ and $x \prec_b z \prec_b y$ and $\alpha \prec_b \gamma \prec_b \beta$ and game tree:

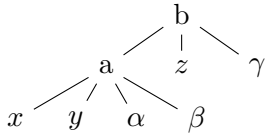

The preferences avoid the forbidden pattern from Theorem 31 (1.); but the pattern is present in any linear extension. In the Nash equilibria of the game, player b is choosing either z or γ; and player a is choosing x or α. In particular, the potential equilibrium outcomes are z and γ – precisely those outcomes that are not weakly Pareto-optimal (because every player prefers y to z and β to γ). Both y and β would even have been Pareto-optimal.

Example 33. We consider a finite two-player game with outcomes $\{x, y, z, t\}$, preferences $t, z \prec_a x, y$ and $x \prec_b z \prec_b y \prec_b t$ and game tree:

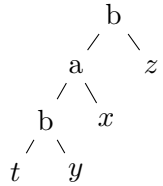

The preferences are strict weak orders and avoid the forbidden pattern from Theorem 31 (1.); but the pattern is present in any linear extension. The only equilibrium outcome is z, despite everyone preferring y ([12]).

7 Games played on finite graphs

A class of infinite sequential games that has received particular interest are games played on finite (directed) graphs. Typical examples are parity games, Muller games, mean-payoff games and discounted payoff games. In each case, the players jointly move a token through the graph.

[12] It may be an interesting remark that in this game every player would benefit, if b could not choose t at his second move.

Each vertex is controlled by some player, and the player controlling the vertex where the token is placed at time t then chooses some outgoing edge to move the token along. Essentially, the finite graph just provides a succinct description of the game tree, which is recovered by fixing a starting vertex and then unraveling the graph.

In the context of games played on finite graphs, often the focus lies on particularly simple strategies, namely those that can be realized by a finite automaton. A two-player win/lose game is called finite-memory determined, if one player has a winning strategy that can be realized by a finite automaton. Note that this strategy has to be able to beat *all* strategies of the opponent, not just those that can also be realized by finite automata.

We will present a transfer result for this setting, showing that finite memory determinacy of the two-player win/lose-games of a class implies the existence of finite memory Nash equilibria for the multi-player multi-outcome case. This result is more general than a similar one obtained by Brihaye, De Pril and Schewe [3],[34, Theorem 4.4.14]. A particular class of games covered by our result but not the previous one are energy parity games as introduced by Chatterjee and Doyen [4].

Theorem 34. Consider a game played by the set A of players on a finite graph such that

1. The $\prec_a$ are automatic-piecewise (with n pieces) prefix-linear strict weak orders.
2. All one-vs-all threshold games of all future games are determined via strategies of memory size m.

Then the game has a Nash equilibrium in strategies of memory size $|A|(m + \log n + K) + 1$, where 2^K is an upper bound on the size of an automaton deciding the pieces for $\prec_a$.

The proof of our main result will have a very similar structure to that of Theorem 27 (which can easily be adapted to finite memory strategies). However, we will need new lemmata to take the roles of Lemma 20 and Lemma 21, as these do not translate to the finite memory setting. The notion of the guarantee also requires some adaption.

Definition 35 (Player (best) guarantee)**.** Let $g = \langle (V, E), v_0, A, \{V_a\}_{a\in A}, (\prec_a)_{a\in A}\rangle$ be a game where $\prec_a$ is a strict weak order for some $a \in A$. For all $h \in H$ ending with some v_1 and strategies s_a for a in g_h let $\gamma_a(h, s_a) := \{v_1 p \in [H_{g_h}] \mid \exists v_1 p' \in [H_{g_h}(s_a)], \neg(hp \prec_a hp')\}$ be the player guarantee by s_a at h. Let $\Gamma_a(h) := \bigcap_{s_a} \gamma_a(h, s_a)$ be the best guarantee of a at h.

For the following lemma, note that whenever the histories h and h' both end in the same vertex v_1, then any strategy s_a for a in g_h can also be understood as a strategy in $g_{h'}$.

Lemma 36. Let g be a game on a finite graph with automatic-piecewise prefix-linear strict weak order $\prec_a$ for some $a \in A$. Let $h \in H$ and $h' \in \overline{h}$. Then $\gamma_a(h, s_a) = \gamma_a(h', s_a)$ for all strategies s_a for a in g_h, and $\Gamma_a(h) = \Gamma_a(h')$.

Proof. Let $v_1 p \in \gamma_a(h, s_a)$, so there exists $v_1 p' \in [H_{g_h}]$ such that $\neg(hp \prec_a hp')$, so $\neg(h'p \prec_a h'p')$ by automatic-piecewise prefix linearity, and $v_1 p \in \gamma_a(h', s_a)$. It shows that $\gamma_a(h, s_a) = \gamma_a(h', s_a)$, which in turn implies $\Gamma_a(h) = \Gamma_a(h')$ by intersection. □

Lemma 37. Let g be a game on a finite graph with strict weak order $\prec_a$ for some $a \in A$. In the threshold games for a of the future games let us assume that if the first player has a winning strategy, she has one with memory size $m \in \mathbb{N}$. Then for all $h \in H$ there exists a strategy s_a with memory size m such that $\gamma_a(h, s_a) = \Gamma_a(h)$.

Proof. For all $p \in [H_{g_h}]$, we have $p \notin \Gamma_a(h)$ iff the first player has a winning strategy in the threshold game for a and p in g_h. In this case let s_a^p be a winning strategy with memory size m. For a game with n vertices there are at most $(n2^m)^{(n2^m)}$ strategy profiles using m bits of memory (by the σ representation), so the s_a^p are finitely many, so at least one of them, which we name s_a, wins the threshold games for all $p \notin \Gamma_a(h)$. This shows that $\gamma_a(h, s_a) \subseteq \Gamma_a(h)$, so equality holds. □

Lemma 38. Let g be a game on a finite graph, and let some $\prec_a$ be an automatic-piecewise prefix-linear strict weak order with n classes. Let K bits be the memory required to decide the pieces for $\prec_a$. If for all $h \in H$ there is a strategy s_a^h with memory size m such that $\Gamma_a(h) = \gamma_a(h, s_a^h)$, there is a strategy s_a in g using no more than $m + \log n + K$ bits and such that $\gamma_a(h, h' \mapsto s_a(hh')) = \Gamma_a(h)$ for all $h \in H(s_a)$.

Proof. By Lemma 36 we can choose our witnessing finite memory strategies s_a^h such that whenever $h' \in \overline{h}$ then $s_a^h = s_a^{h'}$.

We construct a strategy for player a that does two things in parallel: it keeps track of the piece the current history is in (using K bits memory), and it simulates one of the n strategies s_a^h (using m bits for the strategy, and $\log n$ bits to remember which one it currently simulates) until it becomes non-optimal. More specifically, let a start by following the strategy $s_a^{v_0}$. If the run ever reaches a history $h \in H(s_a^{v_0})$ such that $\gamma_a(h, h' \mapsto s_a^{v_0}(hh')) \neq \Gamma_a(h)$, let a start following s_a^h (with strategy-specific memory set to 0^m). This local optimization of the initial strategy $t_0 := s_a^{v_0}$ into t_1 (still a strategy for a in g) does not change the initial guarantee, *i.e.* $\gamma_a(v_0, t_1) = \Gamma_a(v_0)$. Let a repeat this construction as long as necessary.

We now argue that local optimizations occur at most $n - 1$ times along any infinite run. Indeed let us assume that the current run is hh' where $hh' \in \overline{h}$, let t_k and t_{k+p} be a's strategies when leaving h and when arriving at hh', respectively. By Lemma 36 and construction $\Gamma_a(h) = \gamma_a(h, t_k) = \gamma_a(hh', t_k) = \Gamma_a(hh')$, and $\gamma_a(hh', t_{k+p}) \subseteq \gamma_a(hh', t_k)$, so equality holds. This means that once the current strategy has been locally optimized for a piece of histories, its later modifications will also be optimal for the the same piece of histories. Once optimal, always optimal. It implies that local optimization occurs at most $n - 1$ times along every infinite run. □

Proof of Theorem 34. We combine Lemmas 37 and 38 to obtain a finite memory strategy s_a for each player a realizing their guarantee at each vertex, *i.e.* such that $\gamma_a(h, h' \mapsto s_a(hh')) = \Gamma_a(h)$ for all $h \in H(s_a)$. The play p_{NE} induced by this strategy profile will be the play induced by the Nash equilibrium. First note that for every player a and prefix hv of p_{NE}, the suffix of p_{NE} starting after h belongs to $\Gamma_a(hv)$. Now we need to "punish" anyone who deviates, and it is sufficient to make sure that a deviant a at h gets at best the worst of $\Gamma_a(h)$.

For every finite prefix h of p_{NE} let the last vertex of h belong to some V_a. Let us assume that $\Gamma_a(h)$ has no $\prec_a$-least element, and let $p \in \Gamma_a(h)$, so there exists a $\prec_a$-smaller $p' \in \Gamma_a(h)$. Since a has no winning strategy for the threshold game for p' (of the future game at h), the determinacy assumption implies that the coalition of her opponents has one with memory size m. These strategies are finitely many, so one of them, $s_{\overline{a}}$, wins the threshold game for all $p \in \Gamma_a(h)$. So $H_{g_h}(s_{\overline{a}}) \cap \Gamma_a(h) = \emptyset$, which contradicts Lemma 37. So $\Gamma_a(h)$ has a $\prec_a$-minimal element p. Since a has no winning strategy for this threshold, the opponents have one with memory size m. A priori one would need one coalition winning strategy per history h, but similar to Lemma 36, a strategy that works for h also works for every $h' \in \overline{h}$, so only n such strategies are needed.

The final strategy of player a is as follows: follow s_a (introduced above) until one player deviates, when a player b deviates at h, take part in the "punishing" coalition such that b cannot get more than the $\prec_b$-minimum of $\Gamma_b(h)$, which is a lower bound for p_{NE} as mentioned above. If all players play like this, no one has an incentive to deviate: this constitutes a Nash equilibrium.

Regarding the memory requirements, each player needs to simulate the strategies along the main branch of each other player (in order to be able to detect deviations), this requires $|A|(m + \log n + K)$ bits. One extra bit is used to record whether deviation has occurred. If so, the equivalence class of the current history is known, and m bits of memory suffice for the punishing strategy. This memory can be recycled. □

We shall briefly contrast our Theorem 34 to [34, Theorem 4.4.14] (cf. Brihaye, De Pril and Schewe [3]). First, we use strict weak orders, which are an abstraction of the preferences induced by payoff functions as used in [34]. Second, [34, Theorem 4.4.14] requires either that the preferences are prefix-independent and the induced win/lose games are finite memory determined; or that the preferences are prefix-linear and the induced win/lose games are positionally determined. Thus, e.g. energy parity games are not captured by [34, Theorem 4.4.14], as they are not positionally determined, and the preferences are not prefix-independent.

On the other hand, as a corollary of Theorem 34 we find that multi-player multi-outcome energy parity games have Nash equilibria in finite-memory strategies. Here each vertex v of the game graph is labeled by a priority $i_v \in \mathbb{N}$, and by an energy cost $c_{v,a} \in \mathbb{Z}$ for each player. There are fixed energy capacities $\max_a$ and $\min_a$ per player. The current energy budget b_a at some history h is defined recursively via $b_a(\varepsilon) = 0$ and $b_a(hv) = \max\{\min\{b_a(h) + c_{v,a}, \max_a\}, \min_a\}$. The preference of a player for some infinite play depends on the least priority on a vertex occurring infinitely many times, and the lowest energy budget attained at some finite prefix of the run. We further demand that a greater least energy budget (ever occurring along the run) is preferred to a smaller one.

Observation 39. In multi-player multi-outcome energy parity games

1. the one-vs-all threshold of all future games are finite-memory determined with uniform memory bounds (this is not trivial),
2. the preferences are automatic-piecewise prefix-linear.

Corollary 40. Multi-player multi-outcome energy parity games where the preferences are strict weak orders are finite-memory determined.

Before concluding our section on finite graphs, we shall revisit the pareto-optimal equilibria from Section 6. In a generalized Muller game the preferences depend only on the set of vertices visited infinitely many often (thus, in particular are prefix-independent). The first claim in Theorem 41 below both restates the main theorem in [29] and is a corollary of Theorem 34. The second gives a Pareto-optimality version of it, and is novel.

Theorem 41. Every generalised Muller game on a finite graph has a finite-memory NE. If moreover the preferences satisfy $\forall a, b \in A, \forall x, y, z \in O, \neg(z \prec_a y \prec_a x \wedge x \prec_b z \prec_b y)$, there is even a Pareto-optimal NE.

Proof. Unfolding the game yields an infinite sequential game with finitely many outcomes and a Borel measurable outcome function, so by [19, Theorem 3.2] let s be an NE inducing a play p. If moreover the preferences satisfy $\forall a, b \in A, \forall x, y, z \in O, \neg(z \prec_a y \prec_a x \wedge x \prec_b z \prec_b y)$, by

Theorem 31 let us require that s be Pareto optimal. Let uv be a finite prefix of p such that the vertices in v and the vertices that occur infinitely often in p correspond to the same vertices in the original Muller game. Let the players play uv^ω and threaten any deviant during the prefix uv as in s, and at any subgame/finite path/vertex $uv^n v_1 \dots v_k$ as in $uv_1 \dots v_k$. This defines an NE inducing the same outcome as p but requires only finite memory. □

8 Absence of subgame perfect equilibria

In this section we will show that in the simultaneous absence of continuity and the antagonistic/zero-sum property, even a two-player game with three distinct outcomes may fail to have subgame perfect equilibria. It is a straightforward consequence that moving to ε-subgame perfect equilibria cannot help, either. As our (counter-) Example 16, the valuation function here is Δ^0_2-measurable, hence, in a sense, not very discontinuous. A similar counterexample is also exhibited in [39, Example 3].

Example 42. The following game where $z <_a y <_a x$ and $x <_b z <_b y$ has no subgame perfect equilibrium.

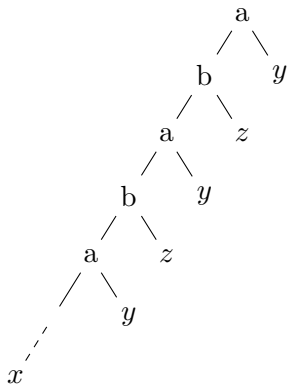

The game above is formally defined as $\langle \{a, b\}, \{0, 1\}, d, O, v, \{<_a, <_b\}\rangle$, where $d^{-1}(a) := 0^{2*}$ and $v(0^\omega) := x$ and $v[0^{2*}1\{0,1\}^\omega] := \{y\}$ and $v[0^{2*+1}1\{0,1\}^\omega] := \{z\}$.

Proof. Assume for a contradiction that there is a subgame perfect equilibrium for this game. Then no subprofile (starting at some node in 0^*) induces the outcome x, because Player b could then switch to the right and obtain z. So for infinitely many nodes in 0^*, Players a or b chooses 1. Also, if Player b chooses 1 at some node 0^{2n+1}, Player a chooses 1 at the node 0^{2n} right above it. This implies that every subprofile rooted at nodes in 0^{2*} induces the outcome y, and subsequently, Player b always chooses 0 at nodes in 0^{2*+1}. But then Player a could always choose 0 and obtain x, contradiction. □

A further example shows us that we can rule out subgame perfect equilibria even with stronger conditions on the functions by using countably many distinct payoffs. This example no longer extends to ε-subgame perfect equilibria.

Example 43. The following game where $y_n := (2^{-n}, 2^{-n})$ and $z_n := (0, 2^{-n-2})$ for all $n \in \mathbb{N}$ has no subgame perfect equilibrium, although the payoff functions are upper-semicontinuous.

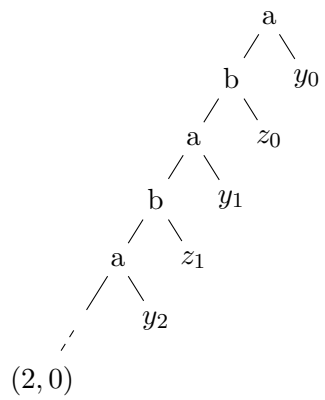

Proof. If the payoffs are $(2,0)$, Player b can improve her payoff as late as required, so there are infinitely many "right" choices in a subgame perfect equilibrium. If the payoffs are not $(2,0)$ then Player a chooses "right" at every node that she owns, so that Player b chooses "left", but then Player a chooses "left" too. □

In [8], an intricate counterexample is provided showing that subgame-perfect equilibria may even fail to exist for probabilistic strategies (in a two-player game with Borel measurable payoff functions).

9 Outlook

There is one open question regarding infinite sequential games with real-valued payoff functions (see the question mark in the Overview table on page 3), namely:

Open Question 44. Do games with countably many players and upper-semicontinuous payoff functions have ε-subgame perfect equilibria?

It seems surprising to have both Nash equilibria and ε-subgame perfect equilibria, but no subgame perfect equilibria – but this is precisely the situation for finitely many players. On the other hand, given the split between finitely many players and countably many players for lower-semicontinuous payoff functions, one should be cautious about assuming that this result would extend. Thus, we do not present a conjecture regarding the answer to the open question.

The results in this paper are generally not constructive – but neither is Nash's theorem in [27], cf. [30]. The extent of non-constructivity is investigated in [23].

Finally, the condition on the payoff functions used in Section 5 seems to merit further investigation. This was that for any sequence $(p^i)_{i\in\mathbb{N}}$ converging to p in C^ω, we find that $\forall i \in \mathbb{N}\ v(p^i) \prec v(p^{i+1})$ implies $\forall i \in \mathbb{N}\ v(p^i) \prec v(p)$. This is a weaker condition than continuity of the function where the upper order topology is used on the codomain, which still seems to be strong enough to formulate some results. In a sense, it is a weakening of continuity that is orthogonal to Borel-measurability. As an example, a result by Gregoriades (reported in [32]) shows that any function of this type from Baire space to the countable ordinals has to be bounded.

References

[1] D. Blackwell (1969): *Infinite G_δ games with imperfect information.* Zastowania Matematyki Applicationes Mathematicae.

[2] Roderick Bloem, Krishnendu Chatterjee, Thomas A. Henzinger & Barbara Jobstmann (2009): *Better Quality in Synthesis through Quantitative Objectives.* In Ahmed Bouajjani & Oded Maler, editors: *Computer Aided Verification*, *Lecture Notes in Computer Science* 5643, Springer Berlin Heidelberg, pp. 140–156, doi:10.1007/978-3-642-02658-4_14.

[3] Thomas Brihaye, Julie De Pril & Sven Schewe (2013): *Multiplayer Cost Games with Simple Nash Equilibria.* In: *Logical Foundations of Computer Science*, LNCS, pp. 59–73, doi:10.1007/978-3-642-35722-0_5. Available at `http://dx.doi.org/10.1007/978-3-642-35722-0_5`.

[4] Krishnendu Chatterjee & Laurent Doyen (2012): *Energy parity games.* Theor. Comput. Sci. 458, pp. 49–60, doi:http://dx.doi.org/10.1016/j.tcs.2012.07.038.

[5] Krishnendu Chatterjee, Thomas A. Henzinger & Marcin Jurdziński (2005): *Mean-Payoff Parity Games.* In: *Proceedings, Twentieth Annual IEEE Symposium on Logic in Computer Science, LICS 2005*, IEEE Computer Society Press, pp. 178–187.

[6] Pieter Collins (2014): *Computable Stochastic Processes.* arXiv:1409.4667.

[7] Martín Escardó (2004): *Synthetic topology of datatypes and classical spaces.* Electronic Notes in Theoretical Computer Science 87.

[8] János Flesch, Jeroen Kuipers, Ayala Mashiah-Yaakovi, Gijs Schoenmakers, Eran Shmaya, Eilon Solan & Koos Vrieze (2014): *Non-existence of subgame-perfect ε-equilibrium in perfect information games with infinite horizon.* International Journal of Game Theory 43(4), pp. 945–951, doi:10.1007/s00182-014-0412-3. Available at `http://dx.doi.org/10.1007/s00182-014-0412-3`.

[9] János Flesch, Jeroen Kuipers, Ayala Mashiah-Yaakovi, Gijs Schoenmakers, Eilon Solan & Koos Vrieze (2010): *Perfect-Information Games with Lower-Semicontinuous Payoffs.* Mathematics of Operations Research 35(4), pp. 742–755.

[10] Drew Fudenberg & David Levine (1983): *Subgame-perfect equilibria of finite- and infinite-horizon games.* Journal of Economic Theory 31(2), pp. 251 – 268, doi:http://dx.doi.org/10.1016/0022-0531(83)90076-5.

[11] D. Gale & F.M. Stewart (1953): *Infinite games with perfect information.* In: *Contributions to the theory of games*, *Annals of Mathematical Studies* 28, Princeton University Press, pp. 245–266.

[12] Erich Grädel & M. Ummels (2008): *Solution concepts and algorithms for Infinite multiplayer game.* In K. Apt & R. van Rooij, editors: *New Perspectives on Games and Interaction*, *Texts in Logic and Games* 4, Amsterdam University Press, pp. 151–178.

[13] Julian Gutierrez & Glynn Winskel (2013): *Borel Determinacy of Concurrent Games.* In: *CONCUR 2013 - Concurrency Theory - 24th International Conference*, pp. 516–530, doi:10.1007/978-3-642-40184-8_36. Available at `http://dx.doi.org/10.1007/978-3-642-40184-8_36`.

[14] Julian Gutierrez & Glynn Winskel (2014): *On the determinacy of concurrent games on event structures with infinite winning sets.* Journal of Computer and System Sciences 80(6), pp. 1119–1137, doi:10.1016/j.jcss.2014.04.005. Available at `http://dx.doi.org/10.1016/j.jcss.2014.04.005`.

[15] Thomas Krieger (2003): *On Pareto equilibria in vector-valued extensive form games.* Mathematical Methods of Operations Research 58, pp. 449–458.

[16] Harold W. Kuhn (1953): *Extensive games and the problem of information.* Contributions to the Theory of Games II.

[17] Stéphane Le Roux (2008): *Generalisation and formalisation in game theory.* Ph.D. thesis, Ecole Normale Supérieure de Lyon.

[18] Stéphane Le Roux (2009): *Acyclic preferences and existence of sequential Nash equilibria: a formal and constructive equivalence*. In: *TPHOLs, International Conference on Theorem Proving in Higher Order Logics*, Lecture Notes in Computer Science, Springer, pp. 293–309.

[19] Stéphane Le Roux (2013): *Infinite Sequential Nash Equilibria*. *Logical Methods in Computer Science* 9(2).

[20] Stéphane Le Roux (2013): *Lazy beats crazy: a spurious yet mathematical justification for laissez-faire*. arXiv:1309.2798.

[21] Stéphane Le Roux (2014): *From winning strategy to Nash equilibrium*. *Mathematical Logic Quarterly* 60(4-5), pp. 354–371, doi:10.1002/malq.201300034. Available at `http://dx.doi.org/10.1002/malq.201300034`. ArXiv 1203.1866.

[22] Stéphane Le Roux & Arno Pauly (2014): *Infinite Sequential Games with Real-valued Payoffs*. In: *CSL-LICS '14*, ACM, pp. 62:1–62:10, doi:10.1145/2603088.2603120. Available at `http://doi.acm.org/10.1145/2603088.2603120`.

[23] Stéphane Le Roux & Arno Pauly (2014): *Weihrauch degrees of finding equilibria in sequential games*. arXiv:1407.5587.

[24] Donald A. Martin (1975): *Borel Determinacy*. *Annals of Mathematics* 102(2), pp. pp. 363–371. Available at `http://www.jstor.org/stable/1971035`.

[25] Donald A. Martin (1998): *The Determinacy of Blackwell Games*. *Journal of Symbolic Logic* 63(4), pp. 1565–1581, doi:10.2307/2586667.

[26] Jean François Mertens (1987): *Repeated Games*. In: *Proc. Internat. Congress Mathematicians*, American Mathematical Society, pp. 1528–1577.

[27] John F. Nash (1950): *Equilibrium Points in N-Person Games*. *Proceedings of the National Academy of Sciences* 36(1), pp. 48–49. Available at `http://www.pnas.org/content/36/1/48.short`.

[28] Rabia Nessah & Kristiaan Kerstens (2008): *Characterizations of the Existence of Nash Equilibria with Non-convex Strategy Sets*. Working Papers 2008-ECO-13, IESEG School of Management. Available at `http://ideas.repec.org/p/ies/wpaper/e200813.html`.

[29] Soumya Paul & Sunil Simon (2009): *Nash Equilibrium in Generalised Muller Games*. In Ravi Kannan & K. Narayan Kumar, editors: *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, *Leibniz International Proceedings in Informatics (LIPIcs)* 4, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, pp. 335–346, doi:http://dx.doi.org/10.4230/LIPIcs.FSTTCS.2009.2330. Available at `http://drops.dagstuhl.de/opus/volltexte/2009/2330`.

[30] Arno Pauly (2010): *How Incomputable is Finding Nash Equilibria?* *Journal of Universal Computer Science* 16(18), pp. 2686–2710, doi:10.3217/jucs-016-18-2686.

[31] Arno Pauly (2012): *On the topological aspects of the theory of represented spaces*. http://arxiv.org/abs/1204.3763.

[32] Arno Pauly (2015): *Computability on the countable ordinals and the Hausdorff-Kuratowski theorem*. arXiv 1501.00386.

[33] Arno Pauly & Matthew de Brecht (2013): *Towards Synthetic Descriptive Set Theory: An instantiation with represented spaces*. http://arxiv.org/abs/1307.1850.

[34] Julie De Pril (2013): *Equilibria in Multiplayer Cost Games*. Ph.D. thesis, Université de Mons.

[35] Rogers Purves & William Sudderth (2011): *Perfect Information Games with Upper Semicontinuous Payoffs*. *Mathematics of Operations Research* 36(3), pp. 468–473.

[36] Tim Roughgarden & Éva Tardos (2007): *Introduction to the Inefficiency of Equilibria*. In: *[?]*, Cambridge University Press, pp. 443–459.

[37] Victor L. Selivanov (2004): *Difference hierarchy in φ-spaces*. *Algebra and Logic* 43(4), pp. 238–248.

[38] Reinhard Selten (1965): *Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit*. *Z. Gesamte Staatswissenschaft* 12, pp. 301–324.

[39] Eilon Solan & Nicolas Vieille (2003): *Deterministic multi-player Dynkin games*. *Journal of Mathematical Economics* 39(8), pp. 911 – 929, doi:http://dx.doi.org/10.1016/S0304-4068(03)00021-1.

[40] Thorsten von Stein (1989): *Vergleich nicht konstruktiv lösbarer Probleme in der Analysis*. Diplomarbeit, Fachbereich Informatik, FernUniversität Hagen.

[41] M. Ummels (2005): *Rational behaviour and strategy construction in infinite multiplayer games*. Diploma thesis, RWTH Aachen.

[42] Klaus Weihrauch (2000): *Computable Analysis*. Springer-Verlag.

Acknowledgements

We are very grateful to Vassilios Gregoriades and the anonymous referees for their helpful comments. This work benefited from the Royal Society International Exchange Grant IE111233. The authors participate in Marie Curie International Research Staff Exchange Scheme *Computable Analysis*, PIRSES-GA-2011- 294962, and are also supported by the ERC inVEST (279499) project.